

# Robust and Rapid Clustering of KPIs for Large-Scale Anomaly Detection

Zhihan Li<sup>1</sup>, Youjian Zhao<sup>1</sup>, Rong Liu<sup>2</sup>, Dan Pei<sup>1</sup>



# Outline

---

- Background
- Algorithm
- Evaluation
- Clustering for KPI Anomaly Detection
- Conclusion



# Outline


---

- Background
- Algorithm
- Evaluation
- Clustering for KPI Anomaly Detection
- Conclusion




# Problem Scenario: KPIs in Internet companies

- Large Internet companies monitor a large number of **KPIs (Key Performance Indicators)**, e.g., CPU utilization, # of queries per second) to ensure the service quality and reliability.
- KPIs are **time series data**. **Anomalies** on KPIs (e.g., a spike or dip) often indicate potential failures on relevant applications, such as server failures, network overload, *etc.*



Idealized time series data




Real-world KPI data (raw)


# Problem Scenario: KPIs in Internet companies

- Large Internet companies monitor a large number of **KPIs (Key Performance Indicators)**, e.g., CPU utilization, # of queries per second) to ensure the service quality and reliability.
- KPIs are **time series data**. **Anomalies** on KPIs (e.g., a spike or dip) often indicate potential failures on relevant applications, such as server failures,

Use **Anomaly Detection** techniques to detect anomalous events timely!



Idealized time series data



Real-world KPI data (raw)

# Problem Scenario: Large-Scale KPI Anomaly Detection

---

- Most anomaly detection algorithms (e.g., Opprentice[1], DONUT[2]) assume that an **individual model** is needed for **each KPI**.
- **Large-scale anomaly detection** is very challenging due to the large overhead of model selection, parameter tuning, model training or anomaly labeling.
- Many KPIs are **similar in underlying shape** due to their implicit associations and similarities.
- Identify **homogeneous KPIs** and apply **one** anomaly detection model per cluster.

# Problem Scenario: Large-Scale KPI Anomaly Detection

---

- Most anomaly detection algorithms (e.g., Opprentice[1], DONUT[2]) assume that an **individual model** is needed for **each KPI**.
- **Large-scale anomaly detection** is very challenging due to the large overhead of model selection, parameter tuning, model training or anomaly

**KPI Clustering** can help!

and similarities.

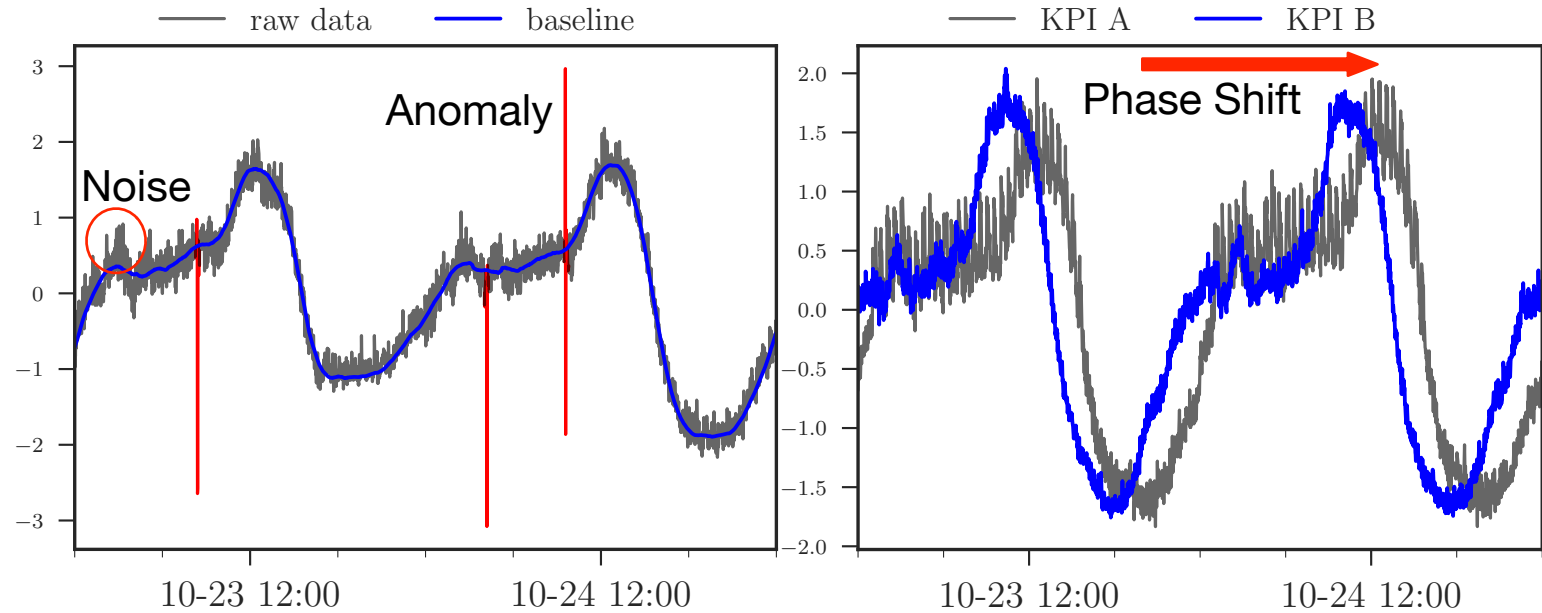
- Identify **homogeneous KPIs** and apply **one** anomaly detection model per cluster.

# Major Challenges

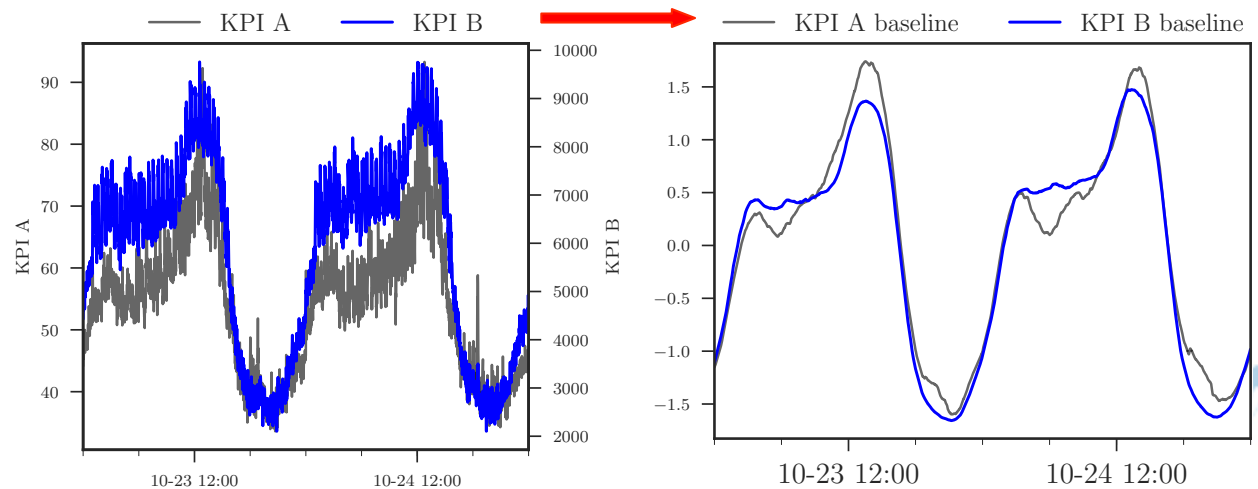
- **Shape Variations**

- Anomalies
- Noises
- Phase Shifts
- Amplitude Differences

- **High Dimensional**



Standardization & Baseline Extraction



Amplitude Difference



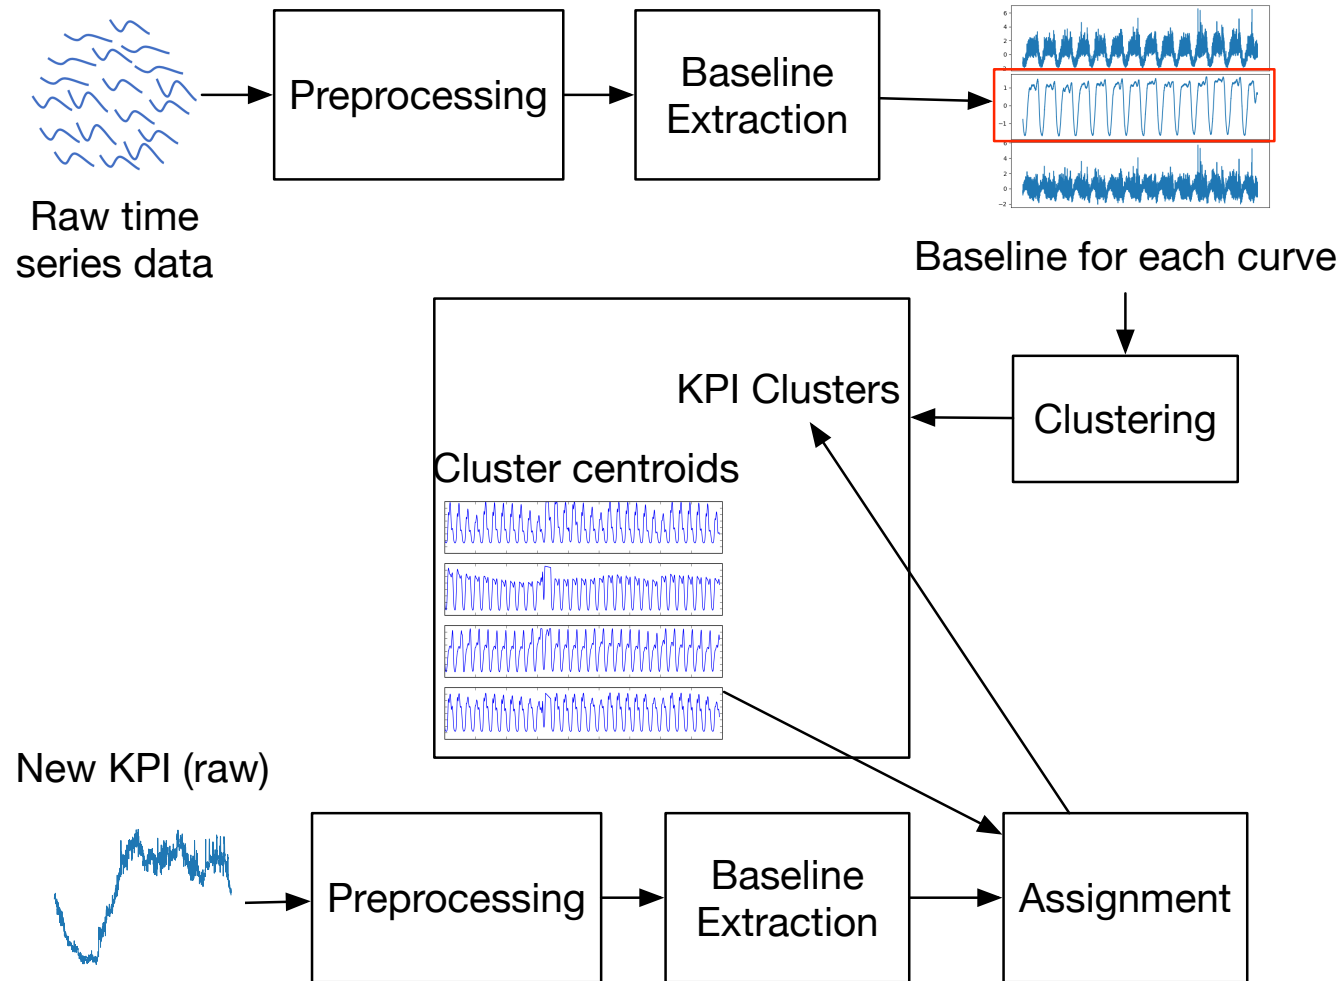
# Outline

---

- Background
- **Algorithm**
- Evaluation
- Clustering for KPI Anomaly Detection
- Conclusion



# Overall Framework of ROCKA



# Preprocessing

---

- Fill missing values with linear interpolation

- **Standardization (remove amplitude differences)**


$$\hat{x}_t = (x_t - \mu_x) / \sigma_x$$

$x_t$  are the original KPI values,  $\mu_x$  and  $\sigma_x$  are the mean and standard deviation of  $x_t$ .



# Baseline Extraction

- Smoothing extreme value
  - Remove the top 5% data which deviates the most from the mean value.
  - Fill them using linear interpolation with their neighboring normal observations.
- Extract baseline
  - Apply moving average with a small sliding window.
  - Baseline extraction removes **anomalies and noises**, while preserving the underlying shape of KPIs.



# Shape-based Similarity Measure

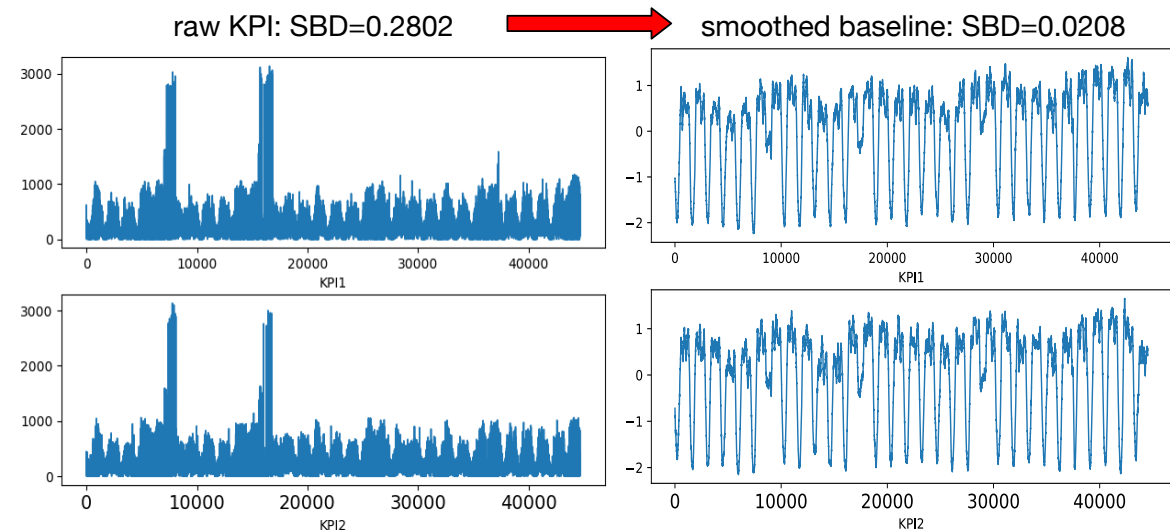
- Normalized version of cross-correlation (NCC)  $\in [-1,1]$ , robust to **phase shifts**.

$$NCC(\vec{x}, \vec{y}) = \max_s \left( \frac{CC_s(\vec{x}, \vec{y})}{\|\vec{x}\|_2 \cdot \|\vec{y}\|_2} \right)$$

- Shape-based distance (**SBD**[3])  $\in [0,2]$ . Smaller SBD means higher shape similarity.

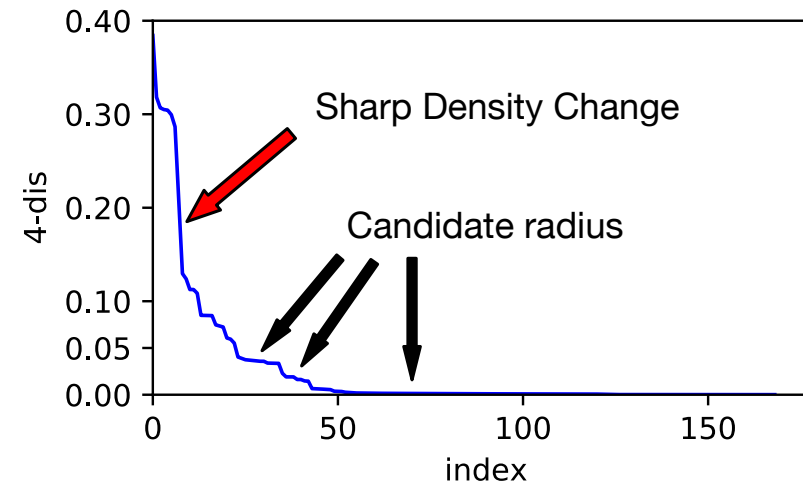
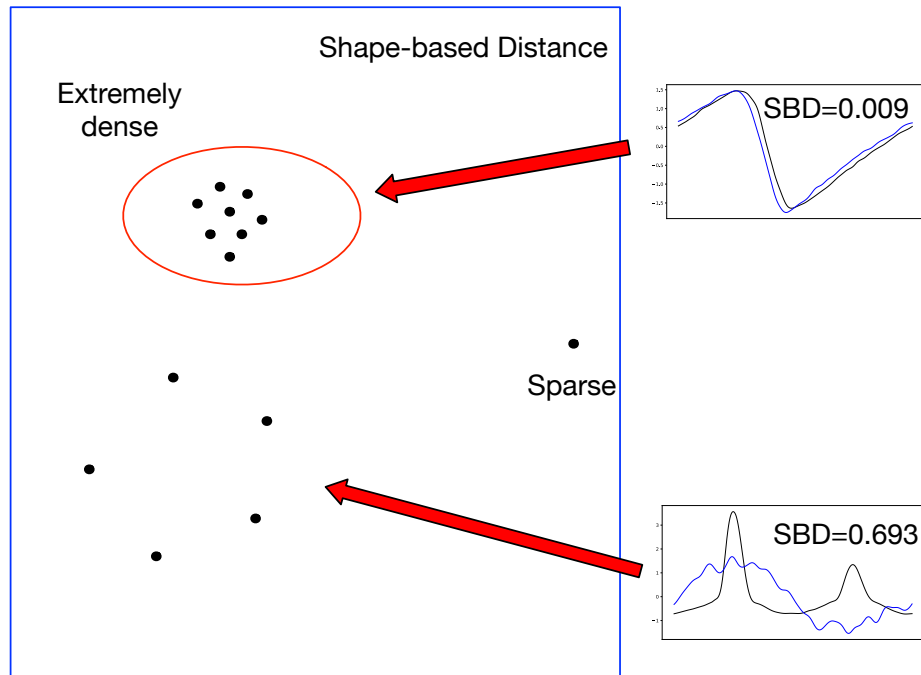
$$SBD(\vec{x}, \vec{y}) = 1 - NCC(\vec{x}, \vec{y})$$

Baseline extraction step plays an important role in finding the shape similarity between KPIs.



# Density-based Clustering

- DBSCAN: find some cores in dense regions, and then expand the cores by transitivity of similarity to form clusters.



With SBD, extremely dense regions contain similar objects that form clusters, while large density radius indicates dissimilar objects.

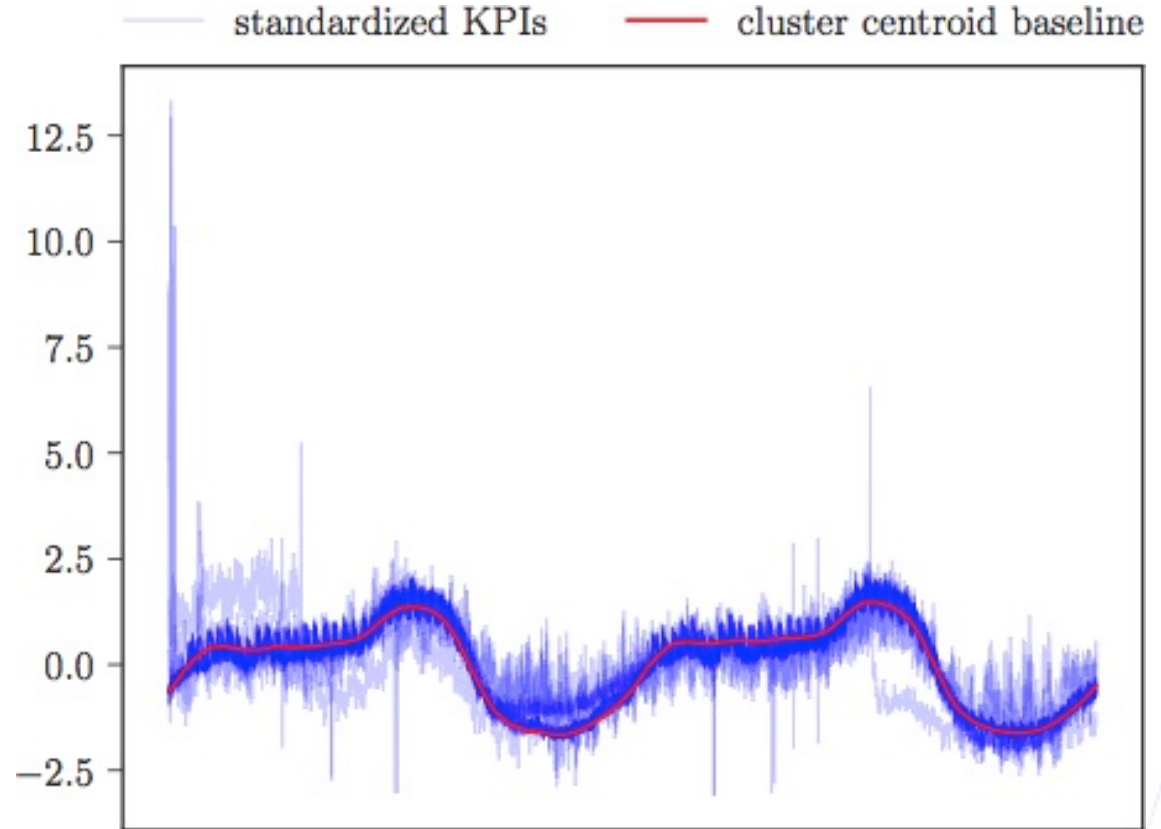
Flat parts on k-dis curve are regarded as candidate radiuses, while steep parts indicate sharp density changes.

# Assignment

- Calculate the centroid of each cluster and assign the rest of KPIs based on centroids.

A cluster with 18 standardized KPIs and its centroid capturing the underlying shape of cluster.

$$\text{centroid} = \arg \min_{\vec{x} \in \text{cluster}_i} \sum_{\vec{y} \in \text{cluster}_i} SBD(\vec{x}, \vec{y})^2$$



# Outline

---


- Background
- Algorithm
- **Evaluation**
- Clustering for KPI Anomaly Detection
- Conclusion





# Performance on public datasets


- YADING[4]: a state-of-the-art clustering algorithm for large-scale time series data.



About 1s for clustering, 0.05s to assign each KPI

# Performance on real-world KPIs

- Evaluation metrics:

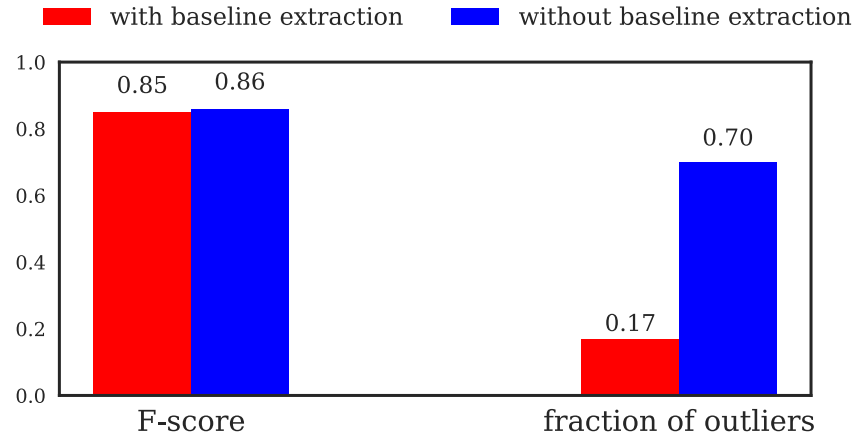


- Each curve is a baseline extracted from the raw KPI

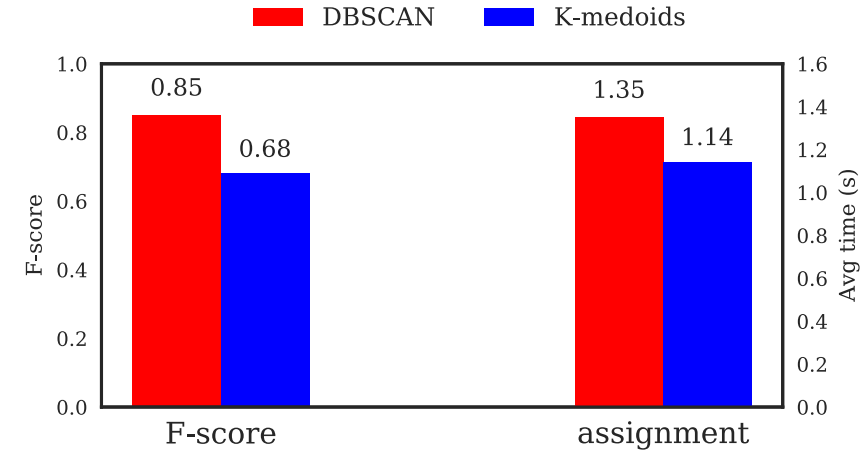
- Performance:

	DS1		DS2	
	<i>ROCKA</i>	<i>YADING'</i>	<i>ROCKA</i>	<i>YADING'</i>
<b>F-score</b>	1.00	0.98	0.85	0.99
<b>fraction of outliers</b>	0.04	0.18	0.17	0.49
<b># clusters</b>	6	7	29	33
<b>avg distance calculation (ms)</b>	53	0.205	58	0.226
<b>avg assignment time (ms)</b>	411	54	1350	99

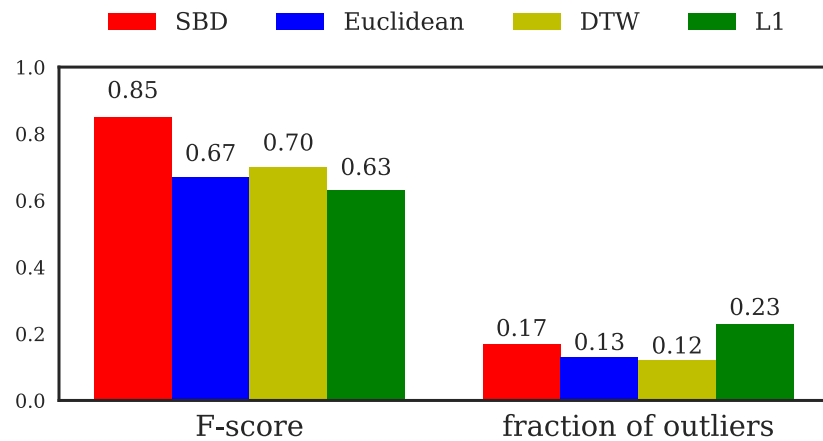
# The effects of techniques in ROCKA



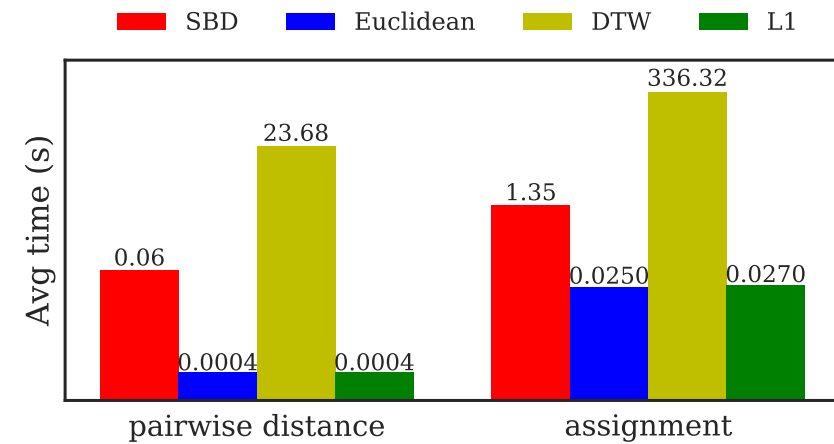
Performance with and without baseline extraction



Performance with different clustering methods



Performance with different similarity measures



Calculation time (log scale) with different similarity measures

# Outline

---

- Background
- Algorithm
- Evaluation
- **Clustering for KPI Anomaly Detection**
- Conclusion



# ROCKA for KPI Anomaly Detection

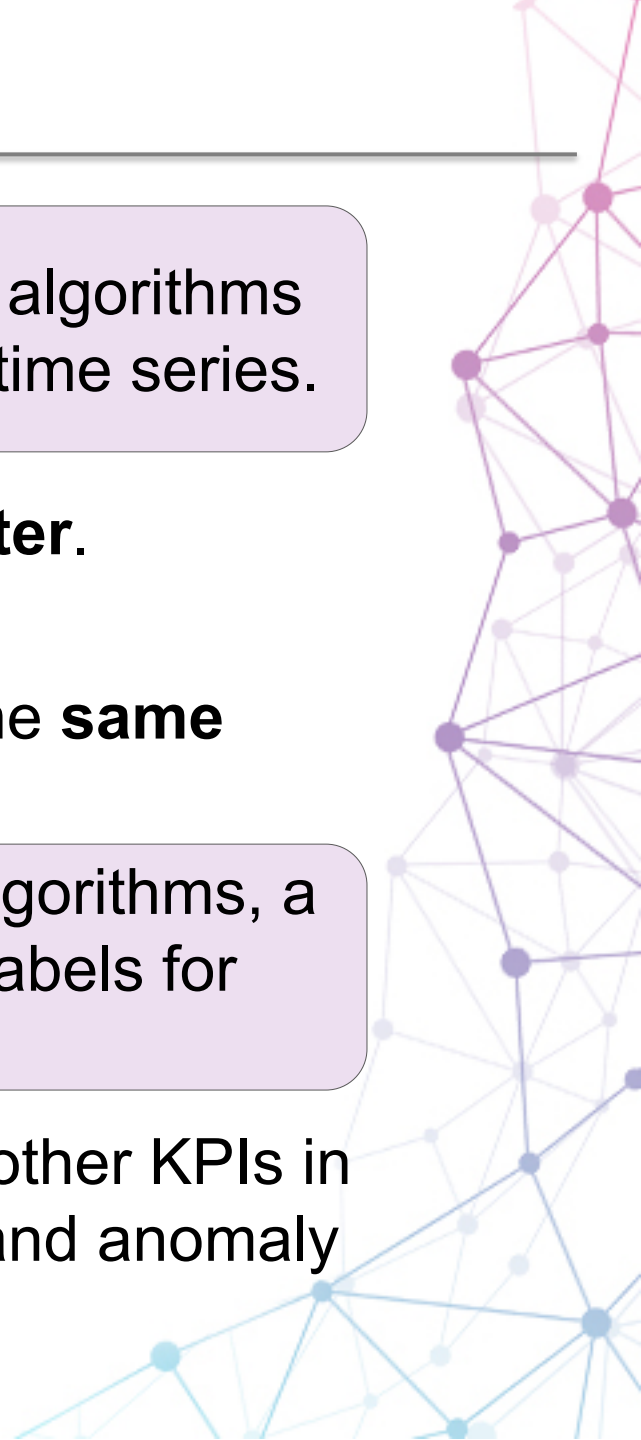
---

Prohibitive amount of **model training time**: Anomaly detection algorithms are often designed to have a model trained for each individual time series.

- ROCKA clusters KPIs similar in underlying shapes into a **cluster**.
- Train anomaly detection model on each cluster **centroid**.
- Directly use the **model** to detect anomalies on other KPIs in the **same cluster**.

Simplifying **threshold selection**: in some anomaly detection algorithms, a threshold needs to be fine-tuned by the ground-truth anomaly labels for optimal performance.

- The **threshold** selected for a cluster centroid can be used by other KPIs in the same cluster, reducing the overhead of parameter tuning and anomaly labeling.



# Anomaly Detection Experiments Setup

---

- **DONUT**[2]: a state-of-the-art unsupervised anomaly detection algorithm for seasonal KPIs.
- **Dataset**: 48 6-month-long KPIs collected from different machines in a large Internet company. Experienced operators has labeled anomalies on these KPIs according to their domain knowledge to provide a ground truth for anomaly detection.
- **Experiments**:
  - E1: **DONUT only**. use DONUT to train an anomaly detection **model for each KPI** and fine-tune the **threshold for each KPI** for the best F-score.
  - E2: **ROCKA + DONUT**. First apply ROCKA on 48 KPIs to form **clusters**, then use DONUT to train an anomaly detection model **only on the centroid KPI** in each cluster, and select the best threshold according to the ground-truth labels **on the centroid**. The model and threshold are then used to detect anomalies in other KPIs of the same cluster.
  - E3: **ROCKA + DONUT + KPI-specific threshold**. Similar to E2, but reestimate the threshold for each KPI, except centroids, according to its ground-truth anomaly labels to get best performance.

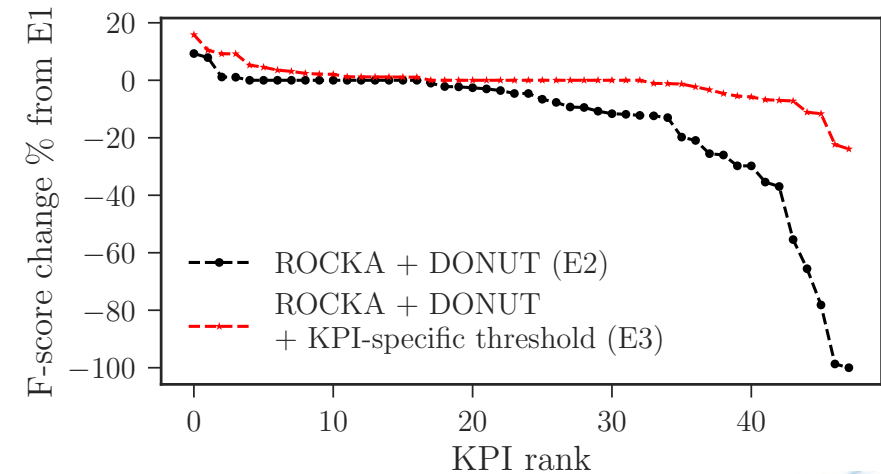
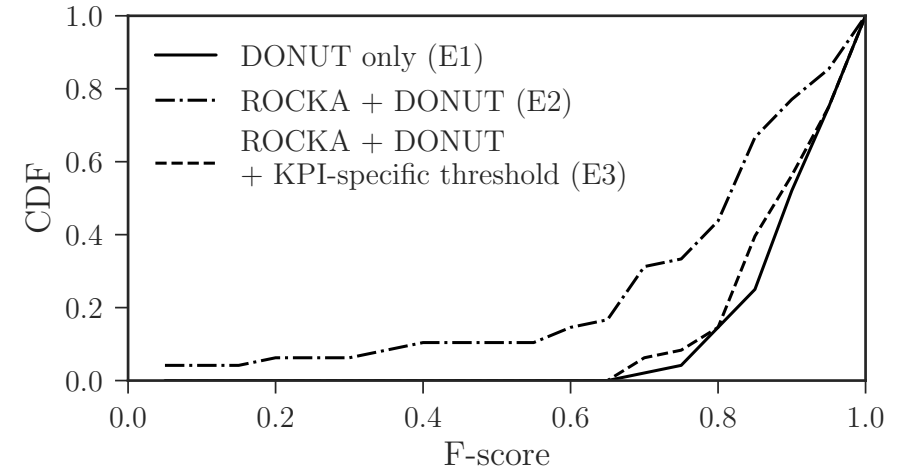
# Anomaly Detection Performance

Cluster	E1	E2	E3	# KPIs
A	0.88	0.66	0.86	18
B	0.79	0.78	0.79	6
C	0.95	0.81	0.95	12
D	0.87	0.86	0.87	4
E	0.90	0.83	0.88	8
Overall	0.89	0.76	0.88	

Average F-score for anomaly detection

algorithm	cluster	tot. train	avg. train	avg. test
DONUT only (E1)	—	51621	1075	345
ROCKA+DONUT (E2)	11	5145	1029	345
ROCKA+DONUT+KPI-specific threshold (E3)	11	5145	1029	345

Time Consumption for anomaly detection (seconds)




The F- score change while using *ROCKA+DONUT*, compared to raw DONUT result (E1)



# Anomaly Detection Performance

Cluster	E1	E2	E3	# KPIs
A	0.88	0.66	0.86	18
B	0.70	0.78	0.70	6

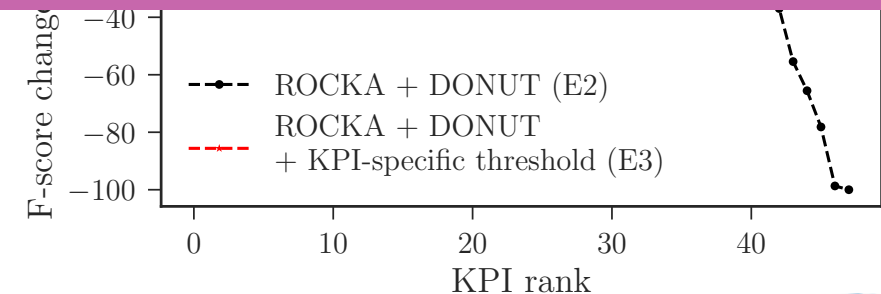


ROCKA reduces the model training time of DONUT by **90%**, with only **15%** performance loss.

When we share model but reestimate the threshold in each cluster, the F-score of most KPIs drop less than **5%** !

<i>ROCKA+DONUT+KPI-specific threshold (E3)</i>	11	5145	1029	345
--	----	------	------	-----

Time Consumption for anomaly detection (seconds)



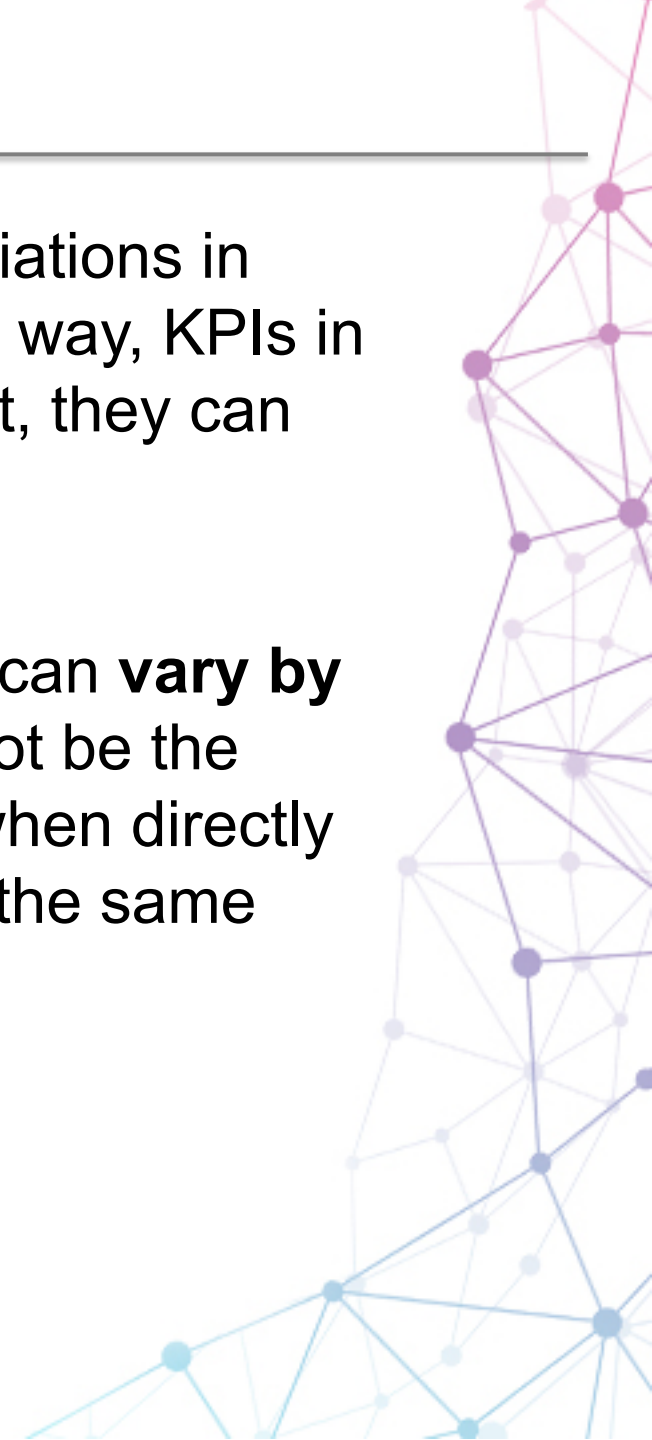
The F- score change while using *ROCKA+DONUT*, compared to raw DONUT result (E1)



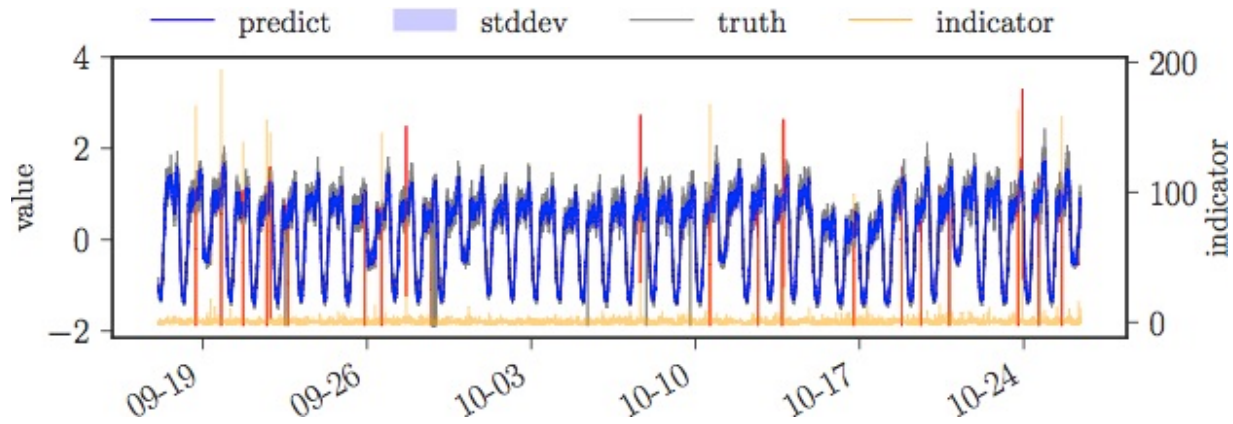
# Analysis of Results

---

- KPIs with similar underlying shape tend to have implicit associations in practice (e.g., belong to the same cluster of machines). In this way, KPIs in the same cluster also have similar normal patterns. As a result, they can **share an anomaly detection model and threshold**.
- KPIs may share the same anomaly detection model, but they can **vary by their anomaly severity levels**, and a uniform threshold cannot be the optimal for every KPI. This leads to some performance drop when directly applying centroid KPI's model and threshold on other KPIs in the same cluster.

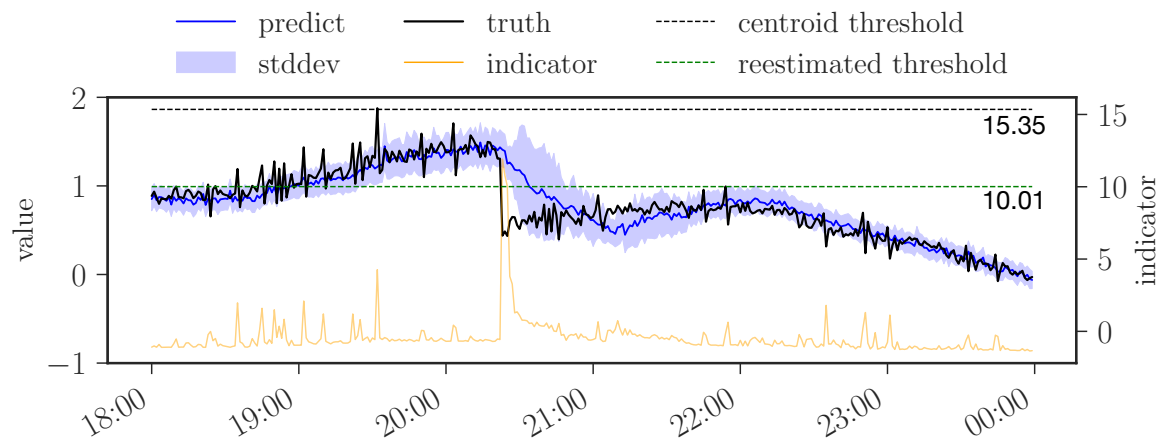


# Analysis of Results



A's cluster centroid model (threshold=15.35)

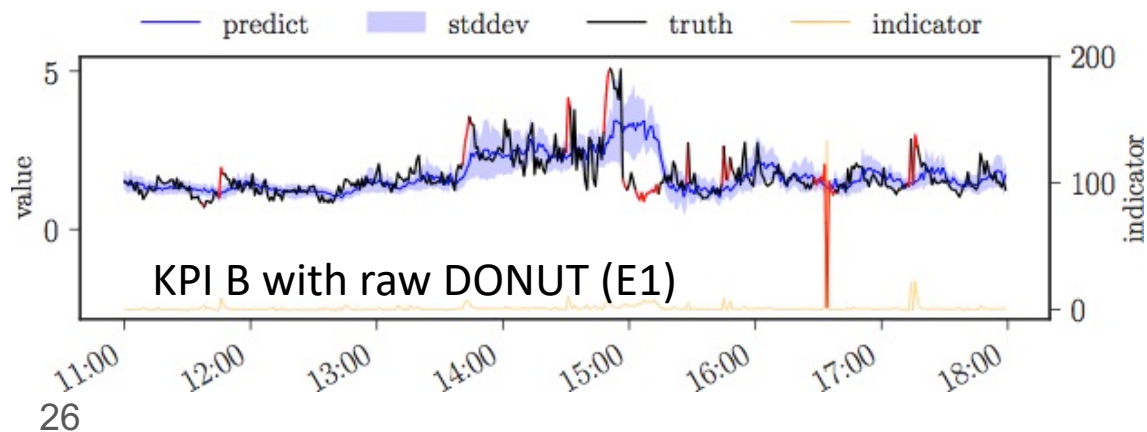
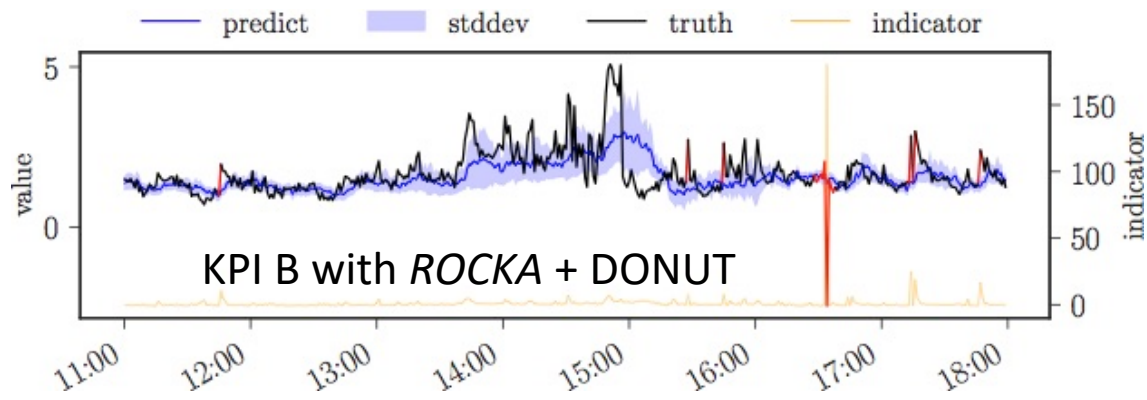
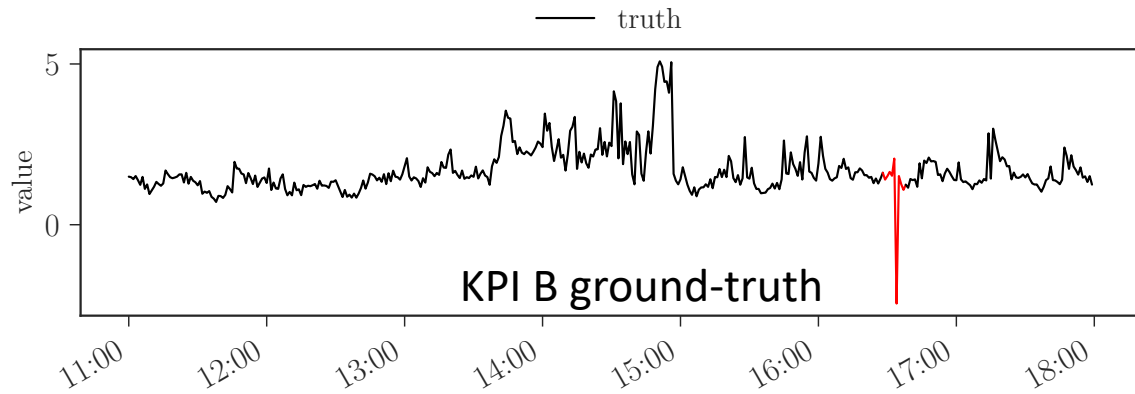
Orange line is anomaly indicator at each point and red line is the anomalies detected by algorithm. The best threshold on KPI A's centroid is 15.35, larger than the indicator of the most significant anomaly on A (11.90). With the reestimated threshold (10.01), all anomalies on A can be detected.



KPI A with ROCKA + DONUT

**anomaly detection model can be shared in the same cluster regardless of different anomaly severity levels.**

# Analysis of Results



The raw DONUT model on KPI B is a bit overfitting and sensitive to small fluctuations. The cluster centroid model is more robust and gets higher F-score.

# Outline

---

- Background
- Algorithm
- Evaluation
- Clustering for KPI Anomaly Detection
- **Conclusion**



# Conclusion

---


- We propose a **robust** and **rapid** time series clustering algorithm, *ROCKA*, to cluster a large number of KPIs, and assist in anomaly detection.
- *ROCKA* **reduces the model training time** of a state-of-the-art anomaly detection algorithm DONUT by **90%**, with only 15% performance loss. This is the first reported study that uses clustering to reduce the training overhead of KPI anomaly detection.
- *ROCKA* is an important first step towards the direction of using KPI clustering to enable large-scale KPI anomaly detection, a key to ensure service reliability in the Internet.

---

# THANK YOU!

Q&A?

[lizhihan17@mails.tsinghua.edu.cn](mailto:lizhihan17@mails.tsinghua.edu.cn)

A decorative network graph pattern is located on the right side of the slide. It consists of a complex web of interconnected nodes and lines, with colors ranging from light blue at the bottom to dark purple at the top. The nodes are represented by small circles, and the lines are thin, creating a mesh-like structure that tapers towards the top right corner.

# References

---

1. D. Liu, Y. Zhao, H. Xu, Y. Sun, D. Pei, J. Luo, X. Jing, and M. Feng, “Opprentice: towards practical and automatic anomaly detection through machine learning,” in *Proc. of IMC*. ACM, 2015, pp. 211–224.
2. H. Xu *et al.*, “Unsupervised anomaly detection via variational auto-encoder for seasonal kpis in web applications,” in *Proc. of the 27th WWW*. ACM, 2018. [Online]. Available: <https://arxiv.org/abs/1802.03903>
3. J. Paparrizos and L. Gravano, “k-shape: Efficient and accurate clustering of time series,” in *Proc. of SIGMOD*. ACM, 2015, pp. 1855–1870.
4. R. Ding, Q. Wang, Y. Dang, Q. Fu, H. Zhang, and D. Zhang, “Yading: Fast clustering of large-scale time series data,” *Proc. of the VLDB Endowment*, vol. 8, no. 5, pp. 473–484, 2015.